State of Iowa
Enterprise IT Standards Program

Enterprise Data Backup Standard
Number: S-TA-011-001

Comments Received during Request for Comment Period
Conducted Aug. 4 – Aug. 18, 2003

August 21, 2003
Compiled by: Doug Kern
Standards Officer
Department of Administrative Services

Enclosed are the comments received during the enterprise-wide Request for Comment Period that ran from August 4 thru August 18, 2003.

A brief summary of the comments is outlined here, followed by the actual comments:

Summary of Comments:

1. Identifying the different solutions to address the two types of data: 1. essential and 2. non-essential, may incur additional costs.

2. For acceptable offsite location storage, what would be the minimum acceptable distance from the operational site? Will the standard guide this decision or will each agency make their own determination?

3. The backup media is to be taken offsite no later than 8 hours of when the backup was taken. Does this 8 hours start at the time the backup is initiated or at the time the backup is completed? Should the backups be offsite prior to the beginning of the next day's business processing?

4. Staffing to perform the handling and offsite deliver of backup media: With automation of backup activities now available, staff may not be onsite on weekends to handle and send tapes to their offsite location. If this activity is required, it would mean that staff would need to be on assignment to perform this work. Do sites that run on Saturday without staff present have to backup on Saturday night?

5. Costs. Additional costs for infrastructure may be needed to accomplish this.

6. There will be a need to distinguish between "system data" and "user data". System data (operating systems, etc.) doesn't typically change as frequently as user data files. Would the backup requirements be different for system data than for user data? In addition, system data backups are usually stored onsite to get quick recovery when needed.

7. Systems and data in remote locations provide an additional challenge. It might be time consuming and inconvenient to retrieve backup tapes on a daily basis at a remote location.

8. Retention periods. What is the retention recommendation?

9. Does “daily” backup mean all 7 days of the week or does it refer to the agency business days (Monday thru Friday)?

10. Does the annual testing requirement mean full recovery of a complete system to a working state? Or just sample recoveries?

Comments Received:.

Comment #1.

From – Judy Peters, Iowa Workforce Development

A. Essential data - agency defined? May be additional cost to identify different technology solutions for essential, non-essential data.

B. "Off-site from its original operational location"
1. Standard does not set a minimum distance from the operational site to the off-site storage location. Standard should either set a minimum distance or state that each agency can set its own distance. If a distance is not set by the standard, the standard should give some guidance for what is acceptable, or industry standards.
2. We use a Data Media vault for tape storage, located in the basement, on a different floor of same building as our data center. Some of our backup tapes are on a weekly rotation from the vault in the basement to the vault at the Lucas Building.

C. "no later than 8 hours of when the backup was taken"
1. Does "taken" mean when the backup begins or ends. Some backup jobs run for 20+ hours, run unattended (no staff on site) over weekends.
2. Use of automated tape libraries have been leveraged to reduce operational staff costs to handle tapes (backups of both essential and non-essential data in same tape library). Staff time required to eject the essential data tapes, package them and move to offsite storage.

D. Cost for compliance with standard.
1. Cost of technology - duplicate and remote equipment for recovery

environment, contract for recovery services - offsite location.
2. Agency Business Continuity Plans and Enterprise Backup/Recovery solutions must balance from feasibility perspective.
----- Forwarded by Judy A. Peters/WA/IWD/EXECUTIVE/IOWA/US on 08/15/2003 10:29 AM -----

---

Comment #2.

From - Pat Clark, ITE

I think you've done a very good job of drafting the backup standard and I have no suggestions for changes. Is there any monitoring process to be defined? Will an agency be asked at any time to demonstrate their ability to perform successful backup and recovery and if so by who?

Comment #3.

From - Chris Bassett, ITE

Comment -

I think that the term "designated data custodian" is pretty clear but I am getting questions about it. So maybe you could define the term.

Comment #4.

From - Doug Lovitt , ITE

STANDARD
<<SNIPPED>>
3. Testing of the backup process will be routinely performed to yield proven successful results
4. The proven ability to retrieve and restore backup data will yield successful results

Comment: I don't understand #4; it looks strikingly similar to #3. It seems to me that if an ability is "proven" it, by definition, yields successful results.

In the same section,

The testing program will accomplish these objectives: 1) demonstrate whether the backup was successful and all intended data was backed up as intended; 2) demonstrate that an

agency can indeed recover the data from their backup media in a timely manner and have the ability to utilize or process the data as may be required; 3) demonstrate that the testing processes and procedures are functional and valid by testing them at least annually.

Comment: in 1), I think the line should read 1) demonstrate that ~~whether~~ the backup...

also, in 2), the word "have" should be "has". (The subject of the verb is "agency.")

also 3) is not really an objective. Do you really want to test the testing process? I think what you want to do is test the backup and recovery process. The objective might be better stated "3) demonstrate the validity of recovery processes and procedures at least annually."

Comment #5.

From - Kay Rozeboom, ITE

You may want to distinguish between system data (file and program definitions, data dictionaries, user ID's, etc.) and user data (people, licenses, income tax filings, etc.).

System data, although critical, does not change as frequently as user data. System data is typically backed up 1-3 times per week, with the most recent generation being kept on site since it is often needed for recovery. So system data backups do not meet the new standard in two ways: 1) backups do not occur daily, 2) most recent backup is not kept offsite. Changing backups procedures for system data to meet the standard would be time-consuming and costly for ITE, with no clear benefit.

Comment #6.

From - Chad Hall, ITE

The only part I have a comment/issue with is the: Backup data is expected to be stored in a location that is offsite from its original operational location as soon as possible and no later than 8 hours of when the backup was taken. part.

I take backups at two locations. One here in DSM and one in Ames. My Ames tapes are in Ames for a week at a time. So, I run a backup and the tapes that are used sit in a tape library until that library is full. Usally1- 3 days. Then the tapes are changed for another fresh set of 13. Then those tapes sit in there for 1-4 days. The tapes that came out of the changer are stored in DOT's vault there in the DOT's computer room. On Thursday, Terri Nelson takes a box of 26 tapes to the DOT. Pulls out the now used 13 tapes and placed them in the box with the other 13 used tapes from 4 days ago and replaces them with 13

fresh tapes from the box she took to Ames. She will then take the used tapes with her to her house then bring them to me on Friday here at ITE. Where I will store the tapes in the locked black cabinets in the back of the server farm.

In short someone would have to drive to and from Ames everyday to adhere to this standard. That's 2 hours minimum a day of time. Terri could pick up the tapes everyday and bring them down. That is time for her every day and what do we do for Friday and Saturday's backup? We could use a currier service but they would have to do the Ames to DSM trip every day. Plus, what would we do about Friday and Saturday's backups?

Also, most of my restore requests are from the Ames tapes after there are already here in DSM. So, sending them to a offsite location would just lengthen the response time to restore a mail box for someone.

The DSM location. I take a backup every night M-F to tape. The job is finished by 8:00pm at night. The next morning I pull the nightly backup and put in a locked cabinet at the back of the server farm and insert the new tape. I then erase the new tape for the next backup. If I had to adhere to the standard some one or me would have to come in and change the tape at about 2:00am. We could get operations staff for this.

Or did I miss the point? Was the backup that goes off site the month end or something??

Comment #7.

From - John Maxwell, ITE

We should also include some recommended retention periods. Right now there is no published standard for legal defense of our practices or for guidelines for the customers.

Comment #8.

From - Larry Grund, Iowa Department of Public Safety

What is the retention objective, and how long is the daily backup media to be retained?

Comment #9.

From - Burley Davis, ITE
(also on behalf of IWD)

First, it seems a reasonable standard and does follow what IWD has outlined in the past. If possible, recommend that the completed standard be incorporated in a request to vendors for a mainframe disaster recovery study. That said, here are some problems and concerns.

Problem and concerns with the proposed standard:

1. Standard does not set a minimum distance from the operational site to the off-site storage location. Standard should either set a minimum distance or state that each agency can set its own distance. If a distance is not set by the standard, the standard should give some guidance for what is acceptable.

2. Does daily really mean daily or at end of business Monday thru Friday? For example, do sites that run lights out on the weekend have to backup Saturday night?

3. The statement that backups should be at off-site storage "no later than 8 hours of when the backup was taken" is too vague. Does "taken" mean when the backup begins or ends. Backing up a terabyte can take hours. Is the time to backup included in the 8 hours? In my opinion, it should be. Ideally, the backup should be off-site before user staff begins next days processing. In many cases, this is not practical. However, if the intent is to recover to the morning of failure, the backup needs to be off-site before 7 am.

4. The testing section of the standard is too vague. Especially in regards to recovery testing. Does the annual recovery testing mean full recovery of a complete system to a working state, or just sample recoveries?

Problems and concerns related to IWD mainframe processing:

1. If the answer to problem 2 above is a Saturday night backup, IWD would need a mainframe operator for several hours Sunday morning to eject, package, and move off-site the Saturday night backups.

2. Daily backups are not possible with the current IWD mainframe technology. It is my opinion that daily backups would require a Shark disk subsystem with the flashcopy feature and 3590 tape drives. If I am right about the equipment, costs would exceed $300,000. Please note: IWD currently does intermediate backups with the HSM product and IWD could do a single weekly backup. However, this method takes longer to recover and the agency has set a recovery time limit of 48 hours from disaster declaration.

3. If the answer to problem 4 above is recovery of a complete system to a working state, IWD cannot currently perform such a test. The last test only used the test LPAR and was unable to test tape library or optical library operation. A full test would require a contract with a hot-site vendor or other state agency that can supply hot-site services.

4. Finally, a brief discussion of what needs to be backed up on the IWD mainframe to meet these requirements. IWD would need disk backups, optical backups, and selected tape backups including all HSM files. One way to accomplish this is by actually doing the backups and moving them off-site. Another method is remote mirroring. To do this, IWD would need to have the following duplicate and remote mainframe equipment: disk subsystem (RVA or SHARK), tape library, and optical library. It is possible to write optical data to tape, so the requirement for a remote optical library could be removed if a larger tape library were used.

Comment #10.

From - John Maxwell, ITE

Requiring people to back up data with no retention guidelines seems pointless to me. I can comply with this standard by making ad hoc copies once in a while when I think of it - in fact that is the process in some places. Without any retention guidelines backing up is moot - you never know if what you want is backed up, where it is stored, or if it can be retrieved.

I agree these are distinct issues, but the main problem is lack of a coherent business continuity plan. Once you establish what is required and how often it changes, you can begin to form some kind of idea of what needs to be backed up, for how long, and on what media.

****

The draft version of the Data Backup standard that was distributed to the enterprise for comment follows:

Enterprise Data Backup Standard
Number: S-TA-011-001

Draft – Version Date 08.01.2003

PURPOSE

To establish a standard to ensure backup copies of electronic data are created so that data availability and retention objectives are satisfied.

INTRODUCTION

Iowa's citizens hold state government to the highest standard for the availability and retention of data. As stewards of citizen information, we are accountable for exercising effective data backup practices, which include frequent back up of state data, secure offsite storage, and reliable data retrieval. To ensure the viability of a data backup and retention program, it is important to systematically test backup and retrieval processes to ensure that these processes and associated technical infrastructure performs in a reliable and accurate manner.

While it is acknowledged that there may be additional costs associated with a backup and retention program to satisfy these expectations, there is also the justification for such activities to meet the accountability and assurance expectations of Iowa's citizens and state government's business programs. This is, in effect, a form of insurance to protect the data of state government, it's citizens, and its stakeholders.

This standard does not attempt to determine which specific data is addressed by the standard. Each agency or data custodian holds the responsibility for making those determinations. Each agency is the best source of knowledge for determining the data they require to accomplish their business objectives and their responsibilities to their stakeholders. It is also understood that determining the data to backup and the length of retention incorporates the application of risk management principles.

In summary, this standard addresses the backup and retrieval of data to ensure data is available in accordance with agency requirements.

STANDARD

This standard consists of 4 major components addressing a data backup and retrieval program:

1. Essential data will be backed up
2. Backup data will be stored at a location that is offsite from its original creation and usage
3. Testing of the backup process will be routinely performed to yield proven successful results
4. The proven ability to retrieve and restore backup data will yield successful results

Each agency or its designated data custodian will backup critical information daily. Daily backups are necessary to help an agency return to a point-in-time state prior to a data loss or a disaster or malfunction to an information system. . Backup data is expected to be stored in a location that is offsite from its original operational location as soon as possible and no later than 8 hours of when the backup was taken.

Each agency or its designated data custodian will be able to retrieve the information from the backup media. Retrieval includes the ability to access, recover, and restore data. A backup is useless if an agency cannot take the information from the backup media and restore it back to production.

Each agency or its designated custodian will be able to demonstrate that their backup, offsite storage, and retrieval processes perform in a consistently reliable and successful manner. Each agency will test their backup and recovery procedures. The testing program will accomplish these objectives: 1) demonstrate whether the backup was successful and all intended data was backed

up as intended; 2) demonstrate that an agency can indeed recover the data from their backup media in a timely manner and have the ability to utilize or process the data as may be required; 3) demonstrate that the testing processes and procedures are functional and valid by testing them at least annually.

The subject data for this standard will be determined by each agency in accordance with agency requirements.